arXiv:1512.00473v2 [physics.chem-ph] 12 Feb 2016

Ab initio molecular dynamics with nuclear quantum effects at classical cost: ring polymer contraction for density functional theory

Ondrej Marsalek[1] and Thomas E. Markland[1, a)]
Department of Chemistry, Stanford University, Stanford, California 94305, USA

(Dated: 15 February 2016)

Path integral molecular dynamics simulations, combined with an *ab initio* evaluation of interactions using electronic structure theory, incorporate the quantum mechanical nature of both the electrons and nuclei, which are essential to accurately describe systems containing light nuclei. However, path integral simulations have traditionally required a computational cost around two orders of magnitude greater than treating the nuclei classically, making them prohibitively costly for most applications. Here we show that the cost of path integral simulations can be dramatically reduced by extending our ring polymer contraction approach to *ab initio* molecular dynamics simulations. By using density functional tight binding as a reference system, we show that our ring polymer contraction scheme gives rapid and systematic convergence to the full path integral density functional theory result. We demonstrate the efficiency of this approach in *ab initio* simulations of liquid water and the reactive protonated and deprotonated water dimer systems. We find that the vast majority of the nuclear quantum effects are accurately captured using contraction to just the ring polymer centroid, which requires the same number of density functional theory calculations as a classical simulation. Combined with a multiple time step scheme using the same reference system, which allows the time step to be increased, this approach is as fast as a typical classical *ab initio* molecular dynamics simulation and 35× faster than a full path integral calculation, while still exactly including the quantum sampling of nuclei. This development thus offers a route to routinely include nuclear quantum effects in *ab initio* molecular dynamics simulations at negligible computational cost.

Keywords: Path integral quantum mechanics, path integral molecular dynamics, multiple time scale molecular dynamics, ring polymer contraction, PIMD, RPMD, CMD, AIMD

I. INTRODUCTION

Obtaining an accurate theoretical description of the atomistic properties of chemical and biological systems requires the development of *ab initio* simulation approaches that explicitly include the quantum mechanical nature of both the electrons and nuclei. *Ab initio* molecular dynamics (AIMD) simulations, where the interactions are obtained from on-the-fly evaluation of the electronic structure, include the quantum nature of the electrons but treat the nuclei as classical particles. Path integral molecular dynamics (PIMD) simulations allow the exact inclusion of nuclear quantum effects (NQEs) in static equilibrium (imaginary time) properties by exploiting the mapping of a quantum system onto an extended classical system of ring polymers comprised of multiple replicas (beads) of the classical system, with harmonic springs that link adjacent replicas of each atom in the system[1–3]. By combining PIMD with AIMD, the resulting *ab initio* path integral molecular dynamics (AI-PIMD) simulations[4–6] thus allow one to capture the full interplay of nuclear and electronic quantum effects.

The imaginary time path integral formalism, which forms the basis of PIMD, also provides the foundation for a number of approximations to the quantum dynamics of chemical systems. These range from methods that utilize imaginary time information to obtain their initial conditions[7], to those which analytically continue the imaginary time data[8,9], to those that utilize the dynamics of the imaginary path itself to approximate quantum dynamics, such as ring polymer molecular dynamics (RPMD)[10,11] and centroid molecular dynamics (CMD)[12].

Despite their utility, PIMD and RPMD have traditionally required a considerable increase in computational cost compared to classical simulations. In particular, the stiff equations of motion result in non-ergodic dynamics, which limits the time step that can be used and also the sampling efficiency[13]. This has led to the development of coordinate transformations and advanced thermostatting methods to alleviate these issues[14–18]. However, even when these issues are addressed, another significant challenge remains: the need to create multiple replicas of the system. Each of these replicas requires the evaluation of the potential energy and forces on all of its atoms. Since obtaining the forces is the rate limiting step in AIMD calculations, the need to create the P replicas makes the cost of an AI-PIMD simulation at least P times higher than the corresponding AIMD simulation. The number of replicas required depends on the highest frequency in the system, and for typical hydrogen-containing systems at room temperature is around 32 for most properties, but can rise as high as 128 if high accuracy of fluctuation properties, such as the heat capacity, is required[19,20].

A number of methods have been proposed to reduce the number of replicas required to converge AI-PIMD simulations. Two particularly popular approaches

a)Electronic mail: tmarkland@stanford.edu

are higher order factorization schemes[21] and approaches based on the generalized Langevin equation (GLE)[17,18]. These methods are able to considerably reduce the number of replicas required, although both suffer from certain limitations. The former requires the use of the Hessian in the calculation of atomic forces, making it impractical for AIMD, where obtaining the Hessian is very expensive. This has led to the introduction of simulation schemes that evolve the system under the standard PIMD Hamiltonian and then utilize thermodynamic reweighting and higher order estimators to recover the rapid convergence with the number of replicas. However, the efficiency of reweighting decreases with system size, limiting its use to small systems[22]. The latter approach of employing colored noise via GLE dynamics has been shown to greatly accelerate the convergence of properties, such as the quantum kinetic energy, in simulations ranging from aqueous systems[23,24] to enzymes[25]. However, although noise matrices have been created that can be applied to a wide range of systems[26], the generation of new ones requires additional parameterization. A given noise matrix will also only accelerate the convergence of the specific correlations of the imaginary time path, such as the virial kinetic energy estimator, that were targeted in the parameterization[18,27]. In addition, due to the non-equilibrium nature of the GLE dynamics, the probability distribution produced cannot be related to a well-defined ensemble, making its combination with Monte Carlo or reweighting unfeasible. Finally, neither of these approaches yield dynamics that can be interpreted in the context of CMD or RPMD, limiting their use to static equilibrium properties.

A particularly appealing scheme that does not suffer from these deficiencies is ring polymer contraction (RPC)[28]. RPC employs a different approach: instead of reducing the total number of replicas, P, it reduces the cost of evaluating the forces on each of them. To achieve this, RPC exploits the fact that the replicas are kept close in space due to the strong harmonic spring terms between them. As a result, any smoothly varying interaction can be approximated with negligible error on a much coarser representation of the imaginary time path, i.e. one with fewer replicas. For systems where the forces can be split into components that vary smoothly in space and those which vary rapidly, one can exploit this observation by evaluating the rapidly varying components on all replicas and the smoothly varying ones on a contracted ring polymer comprised of fewer replicas, P'. If this splitting is constructed such that the computational cost of the rapidly varying forces is negligible compared to that of the smoothly varying forces, one can decrease the cost of the force evaluations by a factor of P/P'. Despite the evaluation of some of the forces on the contracted ring polymer, these forces are exactly projected back onto the full ring polymer. Thus at any instant the positions and forces on the full imaginary time path (ring polymer) representing each particle are known.

The RPC approach has previously been shown to be highly effective in reducing the cost of PIMD and RPMD simulations with empirical potentials by using separation of the inter- and intramolecular forces[28], by range separation of the Coulomb potential[29] and in polarizable force fields by splitting of the contributions to the polarization[30]. In these applications, the properties of interest converged systematically with the number of contracted replicas used and convergence was typically achieved with P'=1 to P'=6, allowing significant increases in efficiency and achieving near classical computational cost with P'=1. In addition, RPC has a well defined Hamiltonian which can be used to assess integration accuracy and allows combination of RPC with reweighting or replica exchange techniques. Finally, since RPC is simply an approximation to some of the interactions, it can be used to perform CMD and RPMD simulations from which approximate quantum dynamics can be obtained. However, despite the success of these splittings for empirical force fields, the interactions obtained from electronic structure methods can not be trivially split into different types.

Here we show it is possible to extend RPC to AI-PIMD while retaining all its desirable features. In particular, we demonstrate that our *ab initio* RPC (AI-RPC) scheme gives rapid convergence to the full AI-PIMD result obtained from density functional theory (DFT) for systems ranging from the reactive protonated and deprotonated water dimers to liquid water, by using self consistent charge density functional tight binding (SCC-DFTB)[31] as a reference system. We show that this convergence can usually be obtained with $P' = 1$ contracted replicas for many properties, with full convergence by $P' = 6$. Further, by exploiting the smoothness in real as well as imaginary time generated by this choice of reference potential, we can increase the time step to 2 fs by using a multiple time scale (MTS) propagation scheme. This combination of MTS and AI-RPC enables a further decrease of the computational cost of the method, yielding a scheme which is in many cases cheaper than a standard classical AIMD simulation while achieving an exact treatment of the NQEs. This development thus opens the door to performing *ab initio* simulations that include NQEs routinely at negligible extra computational cost.

II. THEORY

In this section we briefly review the RPC formalism and outline the details necessary for its combination with *ab initio* evaluation of the interactions by using a reference force.

A. Ring Polymer Contraction

The Hamiltonian of a system of N classical particles of masses m_i is given by

$$H = \sum_{i=1}^{N} \frac{\mathbf{p}_i^2}{2m_i} + V(\mathbf{r}_1, \ldots, \mathbf{r}_N), \tag{1}$$

where in AIMD the potential energy $V(\mathbf{r}_1, \ldots, \mathbf{r}_N)$ is obtained from an electronic structure calculation. The dynamics generated by this Hamiltonian gives rise to classical nuclear sampling of the electronic surface. For distinguishable particles the path integral expression for the quantum partition function corresponding to this Hamiltonian is

$$Z_P = \frac{1}{(2\pi\hbar)^f} \left(\prod_{i=1}^{N} \sqrt{\frac{m_i}{\tilde{m}_i}} \right) \int \mathrm{d}^f \mathbf{p} \int \mathrm{d}^f \mathbf{r}\; e^{-\frac{\beta}{P} H_P(\mathbf{p},\mathbf{r})}, \tag{2}$$

where $f = 3NP$, P is the number of replicas and $\beta = 1/(k_\mathrm{B}T)$. $H_P(\mathbf{p},\mathbf{r})$ is the PIMD Hamiltonian,

$$H_P(\mathbf{p},\mathbf{r}) = H_0(\mathbf{p},\mathbf{r}) + \sum_{j=1}^{P} V(\mathbf{r}_1^{(j)}, \ldots, \mathbf{r}_N^{(j)}), \tag{3}$$

where

$$H_0(\mathbf{p},\mathbf{r}) = \sum_{i=1}^{N} \sum_{j=1}^{P} \left(\frac{|\mathbf{p}_i^{(j)}|^2}{2\tilde{m}_i} + \frac{1}{2} m_i \omega_P^2 |\mathbf{r}_i^{(j)} - \mathbf{r}_i^{(j-1)}|^2 \right) \tag{4}$$

is the free ring polymer Hamiltonian, with cyclic boundary conditions, $j+P \equiv j$, implied. The dynamical (sampling) masses $\tilde{m}_i$ do not have to be the same as the physical masses m_i and in PIMD simulations are typically set based on computational convenience. We use the notation

$$\mathbf{p} \equiv \{\mathbf{p}_i^{(j)}\}_{i=1\ldots N}^{j=1\ldots P},\; \mathbf{r} \equiv \{\mathbf{r}_i^{(j)}\}_{i=1\ldots N}^{j=1\ldots P} \tag{5}$$

for the full set of ring polymer momenta and positions, respectively. The Hamiltonian in Eq. 3 corresponds to a set of P copies of the physical system, where the adjacent replicas representing each particle are connected by harmonic springs of frequency $\omega_P = P/(\beta\hbar)$. As $P \to \infty$, the classical dynamics generated by this Hamiltonian samples the exact quantum mechanical partition function. However, evaluating this Hamiltonian requires P electronic structure calculations, to compute $V(\mathbf{r}^{(j)})$ for each of the P replicas.

Let us now consider splitting the full force on a particle $\mathbf{f}_\mathrm{full}$, which requires a full electronic structure calculation to obtain, into a reference component $\mathbf{f}_\mathrm{ref}$ and the remaining difference force,

$$\mathbf{f}_\mathrm{diff} = \mathbf{f}_\mathrm{full} - \mathbf{f}_\mathrm{ref}. \tag{6}$$

If one constructs the reference system such that it captures the rapidly varying parts of the full interactions, the remaining difference force will be smoothly varying in space. If this is the case, the smoothly varying force can be accurately approximated by evaluating it on a contracted ring polymer. This is achieved by mapping the P-replica ring polymer positions, $\mathbf{r}$, onto a contracted set of positions, $\mathbf{r}' \equiv \{\mathbf{r}_i^{(j')}\}_{i=1\ldots N}^{j'=1\ldots P'}$, of $P' \leq P$ replicas using the transformation

$$\mathbf{r}_i^{(j')} = \sum_{j=1}^{P} T_{j'j}\, \mathbf{r}_i^{(j)}. \tag{7}$$

A number of ways to achieve this transformation are possible. Here we use our previously introduced normal mode contraction procedure, where the ring polymer is transformed to its local normal mode representation. The $P - P'$ normal modes with the highest frequencies, which interact most weakly with the physical potential, are then discarded and the remaining P' normal modes are transformed back to the coordinate representation[28]. The required transformation matrix, $T_{j'j}$, is given in Ref. 28 and can be simply evaluated and applied at negligible computational cost. In the limiting cases, the net effect of this transformation is such that when $P' = P$, it leaves the ring polymer unchanged and when $P' = 1$ it contracts the ring polymer to its centroid, $\bar{\mathbf{r}}_i$,

$$\mathbf{r}_i^{(j'=P'=1)} = \bar{\mathbf{r}}_i = \frac{1}{P} \sum_{j=1}^{P} \mathbf{r}_i^{(j)}. \tag{8}$$

For intermediate values, the transformation creates a contracted set of P' positions which approximately represent the full ring polymer (effectively taking a lower-order Fourier representation of the imaginary time path). The slowly varying difference force can then be evaluated on the contracted positions, $\mathbf{f}_{i,\mathrm{diff}}^{(j')}$, and then projected back onto the full ring polymer representing each particle:

$$\mathbf{f}_{i,\mathrm{diff}}^{(j)} = \frac{P}{P'} \sum_{j'=1}^{P'} T_{j'j} \mathbf{f}_{i,\mathrm{diff}}^{(j')}, \tag{9}$$

in which the factor of P/P' on the right-hand side arises naturally from the ratio of the number of replicas in the full and contracted system[28]. With this contracted approximation to the difference force, the full force on each ring polymer is

$$\begin{aligned} \mathbf{f}_i^{(j)} &= \mathbf{f}_{i,\mathrm{ref}}^{(j)} + \frac{P}{P'} \sum_{j'=1}^{P'} T_{j'j} \mathbf{f}_{i,\mathrm{diff}}^{(j')} \\ &= \mathbf{f}_{i,\mathrm{ref}}^{(j)} + \frac{P}{P'} \sum_{j'=1}^{P'} T_{j'j} \left(\mathbf{f}_{i,\mathrm{full}}^{(j')} - \mathbf{f}_{i,\mathrm{ref}}^{(j')} \right), \end{aligned} \tag{10}$$

where the second equality follows from the definition of the difference force (Eq. 6). Hence the computationally expensive full force only needs to be evaluated P' times for each configuration, rather than P times as in a standard PIMD simulation.

Although the forces in Eq. 10 are an approximation to the full forces on each replica, the dynamics generated from them formally conserve a well defined Hamiltonian[28]

$$H'_P(\mathbf{p},\mathbf{r}) = H_0(\mathbf{p},\mathbf{r}) + \sum_{j=1}^{P} V_{\text{ref}}^{(j)}(\mathbf{r}) + \frac{P}{P'}\sum_{j'=1}^{P'} V_{\text{diff}}^{(j')}(\mathbf{r}'), \tag{11}$$

where V_{ref} and V_{diff} are the potential energies that correspond to the reference and difference forces, respectively, and together replace the potential energy term in Eq. 3. The existence of this Hamiltonian allows energy conservation to be checked during dynamics to assess integration accuracy and also enables the combination of RPC with reweighting schemes and methods that require a well defined ensemble.

It is important to note that since the contracted forces are projected back onto the full ring polymers which are then evolved in time under these forces, the positions of all the replicas that comprise full ring polymer are known. These P replica positions can therefore be used to evaluate any position dependent property at each time step. For both forces and potential energies, one naturally generates the contracted estimators for them during the RPC evolution (Eqs. 10 and 11). The forces on each replica obtained in a PIMD simulation can be used to compute the quantum kinetic energy of a particle, i, as the ensemble average of the virial estimator[32],

$$\mathcal{T}_i(\mathbf{r}) = \frac{3k_{\text{B}}T}{2} - \frac{1}{2P}\sum_{j=1}^{P}(\mathbf{r}_i^{(j)} - \bar{\mathbf{r}}_i)\cdot\mathbf{f}_i^{(j)}. \tag{12}$$

However, for these properties an alternative approach to achieve more accurate results from a RPC trajectory is to use uncontracted estimators (UE)[28]. These UEs are evaluated by performing full (uncontracted) potential and force evaluations on the full P replica configurations generated from the contracted simulation. This post-processing of the RPC configurations requires extra full potential and force evaluations and hence some additional cost. However, by choosing to only post-process a subset of frames, one can take advantage of the fact that adjacent time steps are heavily correlated, and hence evaluating the UE adds comparatively little computational overhead.

B. Constructing a reference system

RPC provides a way to save computational effort by using a physically motivated approximation in the evaluation of physical interactions in the system while still retaining the full dimensionality of the P-replica imaginary time path. The form of our contraction transformation, Eq. 7, and the definition of the difference force, Eq. 6, guarantee convergence as $P' \to P$. However, the rate of this convergence and the computational cost savings obtained from reducing the number of replicas in the contracted ring polymer depend crucially on the creation of an efficient reference system that accurately captures the rapidly varying parts of the interactions in the full system and thus leaves a slowly varying difference force.

The objective is therefore to construct a reference force f_{ref} which, when taken from the full electronic structure force f_{full}, leaves a remaining force f_{diff} that can be well approximated on the contracted polymer. The reference force for an AI-RPC approach should thus satisfy the following requirements:

1. It is computationally quick to evaluate compared to performing the full electronic structure calculation (ideally by at least a factor of 10).

2. It gives a difference force that is smoothly varying in space.

3. It does not make assumptions on chemical bonding, i.e. it allows bond breaking and formation as the simulation progresses.

It is vital to note that the reference force only has to leave a slowly varying remainder — i.e. the reference force can be something that would give a very poor result for the dynamics and structure of the system if used alone (without the difference force which corrects for its deficiencies). For example, when applied to empirical potentials, RPC has previously been shown to be successful even when the reference force was chosen as just the intramolecular force[28]. Hence, used alone, this reference system would predict no interactions between molecules (an ideal gas). However, the presence of the difference force on the contracted replicas corrects for this, and it was shown[28] that convergence of many structural properties of liquid water was obtained even with $P' = 1$.

The idea of a reference system has natural origins in the MTS molecular dynamics community, where the reversible reference systems propagator algorithm (r-RESPA)[33] was formulated as a method to allow efficient time-reversible symplectic propagation in molecular dynamics simulations with multiple components of the force which vary on different time scales. Whereas MTS schemes exploit the slowly varying nature of some forces in real time to extend the propagation time step, RPC takes advantage of the spatially smooth variation of the forces in the imaginary time of the path integral. Thus, the considerations needed for a good reference force in the two approaches are similar. This suggests that for systems where a suitable reference force can be identified, it can be used to naturally utilize both MTS and RPC in a simulation.

This complementarity is illustrated in Fig. 1, which shows the force autocorrelation function of all the protons in the protonated water dimer along the imaginary time path (left panel) and in real time (right panel) obtained from a 32 replica PIMD simulation. The full DFT force (blue line) decays rapidly along the imaginary time path, which extends from $i\tau = 0$ to $\beta\hbar$ (in the language of the discretized ring polymer, from replica 1 to 32 in

this case). Due to the cyclic boundary conditions on the path, the furthest distance in imaginary time between the replicas is at $i\tau = \beta\hbar/2$ (i.e. the force correlation between replica 1 and replica 16). Likewise, in real time the full force also rapidly decays and additionally exhibits oscillations arising from bond vibrations.

We now consider using SCC-DFTB as the reference force. The green line in Fig. 1 shows the difference force which would be obtained using this choice. In both real and imaginary time, the difference force is roughly 2 orders of magnitude smaller and is also much more smoothly varying than the full force. This suggests that the reference force captures almost all of the variation of the full force both along the ring polymer as well as in real time. Hence the SCC-DFTB reference force is able to provide a smooth and small difference force in both real and imaginary time. In addition, using a SCC-DFTB reference for DFT also satisfies the other desirable features such as the ability to make and break bonds and computational savings of typically 2 orders of magnitude over DFT due to the tabulation of the Hamiltonian and overlap matrix elements and the use of a minimal valence basis set[34].

One could also imagine using a number of other methods that have been suggested to split *ab initio* interactions for MTS schemes to extend RPC to *ab initio* PIMD. For DFT, these range from using a reference force generated by using reduced basis sets and non-iterative Harris functionals[35], neglecting Hartree-Fock exchange in hybrid functionals[36], range seperation of the Coulomb operator[37] or by using a reference empirical force field[37,38]. For higher level wavefunction methods, such as MP2, reference forces that have been suggested include neglecting the dynamic electron correlation[39], aggressive electron integral screening[40], and using DFT as a reference for the higher level theory[41].

C. Multiple time scale integration

Given the natural complementarity of MTS, which exploits the smoothness of the difference force in real time, and RPC, which exploits its smoothness in imaginary time, we combine both approaches in our simulations. Since the utility and theory behind MTS approaches are well established and that we have previously combined RPC and MTS a number of times for empirical potentials[42–45], here we only sketch the relevant details specific to our combination with AI-RPC.

The MTS integration used here was given by the factorization of the propagator

$$\begin{aligned}&\exp\left(i\mathcal{L}\Delta t\right) = \\ &\quad \exp\left(i\mathcal{L}_\gamma \Delta t/2\right)\exp\left(i\mathcal{L}_{\rm diff}\Delta t/2\right)\times \\ &\quad \left[\exp\left(i\mathcal{L}_{\rm ref}\delta t/2\right)\exp\left(i\mathcal{L}_0\delta t\right)\exp\left(i\mathcal{L}_{\rm ref}\delta t/2\right)\right]^M \times \\ &\quad \exp\left(i\mathcal{L}_{\rm diff}\Delta t/2\right)\exp\left(i\mathcal{L}_\gamma \Delta t/2\right),\end{aligned} \tag{13}$$

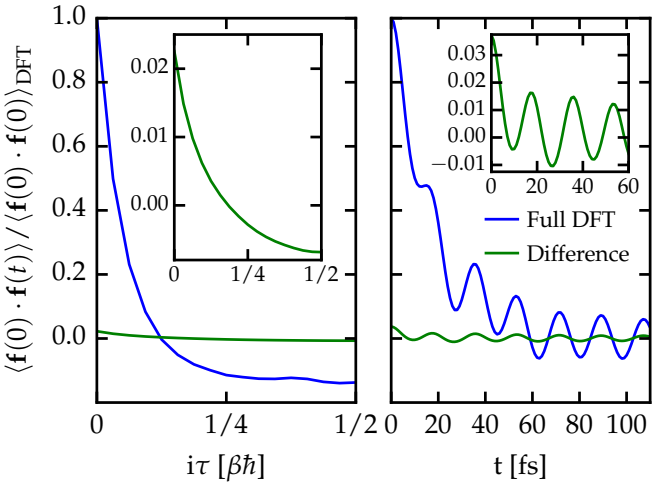

FIG. 1. Force autocorrelation functions along the imaginary time path (left panel) and in real time (right panel) for the protons in the protonated water dimer from a 32 replica PIMD simulation. The correlation function of the full DFT forces is shown in blue, while the correlation function of the difference between the DFT and DFTB forces is shown in green. Correlation functions are scaled relative to the full DFT result at t=0 and $i\tau$=0, $\langle \mathbf{f}(0)\cdot\mathbf{f}(0)\rangle_{\rm DFT}$. The insets show the details of the difference force.

where $\Delta t = M\delta t$ is the outer time step and δt is the inner time step. Here the propagator $\exp\left(i\mathcal{L}_\gamma\Delta t/2\right)$ evolves the system under the thermostat for the time interval $\Delta t/2$ and $\exp\left(i\mathcal{L}_{\rm diff}\Delta t/2\right)$ evolves the system momenta under the difference force by $\Delta t/2$. In the inner loop, $\exp\left(i\mathcal{L}_{\rm ref}\delta t/2\right)$ evolves the system momenta under the reference force by half a small time step $\delta t/2$ and $\exp\left(i\mathcal{L}_0\delta t\right)$ evolves the system under the influence of the free ring polymer Hamiltonian H_0 by transforming to the free ring polymer normal mode representation, which allows analytic integration. The explicit operations of each of these propagators are given in Ref. 16.

One major consideration in MTS schemes, which does not arise in RPC, is the existence of the resonance barrier, which limits the largest outer time step which can be used[46]. The first resonance can be shown to occur when the outer time step exceeds $\Delta t_{max} = \tau/\pi$, where τ is the time-period of the fastest mode in the problem[46]. For liquid water, where the highest physical frequency is the O-H stretch at ∼3600 cm^{-1}, this yields a value of Δt_{max}=2.95 fs. However, the harmonic springs linking the replicas in PIMD give rise to additional high-frequency modes. The highest ring polymer normal mode frequency for the free ring polymer is[16]

$$\omega_{\rm RP,max} = \frac{2P}{\beta\hbar} \tag{14}$$

and for a ring polymer in a physical harmonic potential of frequency ω it shifts to

$$\sqrt{\omega_{\rm RP,max}^2 + \omega^2}. \tag{15}$$

For liquid water with $P = 32$ at T=300 K, the highest frequency in the PIMD simulation is therefore ∼13345 cm^{-1} and hence the first resonance would be expected at Δt=0.8 fs. To allow us to avoid the resonance barrier and thus use larger outer time steps in our MTS simulations, we use the standard procedure of shifting all the normal mode frequencies to a single sufficiently low frequency by adjusting the normal modes' dynamical masses[12,15,47]. In cases where changing the masses is not an option, one could consider coupling the highest normal modes to a strong thermostat to surpass the resonance barrier using either colored noise[48] or targeted white noise[16]. The PILE thermostat scheme achieves the latter by transforming to the normal mode representation and coupling each normal mode to a Langevin thermostat with critical damping based on its free ring polymer frequency. One case where the masses of the normal modes are determined and should not be changed is when obtaining approximate quantum dynamics using the RPMD approach where, $\tilde{m}_i = m_i$. In this case, the TRPMD method[49], which uses the PILE thermostat with a specific choice of damping coefficients to obtain dynamics, might be advantageous, since it also has damping applied to the normal modes of the ring polymer. Indeed, results along these lines do suggest that one may be able to go beyond the first resonance barrier when the PILE equations of motion are used[42,45,50].

III. SIMULATION DETAILS

We performed AI-PIMD and AIMD and AI-RPC simulations of the gas-phase protonated and deprotonated water dimer, and of bulk liquid water at T=300 K. The liquid water simulations were performed under NVT conditions using a system of 64 water molecules with periodic boundary conditions. A cubic box with sides of length 12.42 Å was used, giving a density of 1000.8 kg/m^3.

The benchmark quantum simulation of liquid water used for comparison was performed using the PIGLET scheme[18]. Due to the accelerated convergence with the number of replicas afforded by PIGLET, we use 6 replicas, which has previously been shown to give excellent agreement for the position based properties and quantum kinetic energies to the exact path integral result[18].

In all other benchmark PIMD and in the AI-RPC simulations, the total number of replicas was set to P=32 and the dynamical masses of all ring polymer normal modes were rescaled to shift their frequencies to 500 cm^{-1} to ensure accurate integration of the equations of motion[15]. The inner time step was δt=0.5 fs in all simulations and in trajectories run with MTS, the outer time step was Δt=2.0 fs (M=4). The PILE-G thermostat[16] was used to sample the canonical ensemble. In this thermostatting scheme, a critical damping Langevin thermostat is attached to each ring polymer normal mode of each particle, while the total kinetic energy of the particle centroids is coupled to a single stochastic velocity rescaling thermostat. In our classical MD simulations, the temperature was maintained using a single global stochastic velocity rescaling thermostat[51]. The i-PI program[52], including the implementation of MTS[50], was used to perform all the molecular dynamics simulations. Interactions were obtained by calling external electronic structure programs through the socket interface.

For the gas-phase protonated and deprotonated dimers, DFT interactions were evaluated using the B3LYP hybrid density functional[53] and the 6-311++G** basis set in the Gaussian program[54] called through the calculator provided by the Atomic Simulation Environment[55]. SCC-DFTB[31] interactions including diagonal DFTB3 terms[34] were evaluated in CP2K[56] using the provided parameterization for O and H atoms.

In our liquid water simulations both the DFT and SCC-DFTB interactions were evaluated using the CP2K program[56,57]. For the DFT interactions, we used the revPBE generalized gradient approximation (GGA) density functional[58,59] with the PBE DFT-D3 dispersion correction[60]. Atomic cores were represented using the dual-space Goedecker-Tetter-Hutter pseudopotentials[61]. Kohn-Sham orbitals were expanded in a double ζ contracted Gaussian atomic basis set with polarization functions (DZVP) and a cutoff of 400 Ry was used for the auxiliary plane-wave basis set of the GPW method[62]. The self-consistent field cycle was converged to an electronic gradient tolerance of $\epsilon_{\rm SCF} = 5 \times 10^{-7}$ using the orbital transformation method[63] with the initial guess provided by the always stable predictor-corrector extrapolation method[64,65] at each molecular dynamics step. SCC-DFTB interactions[31], including diagonal DFTB3 terms[34], were combined with the D3 dispersion correction[60] and evaluated in periodic boundary conditions, using Ewald summation for electrostatics.

For the protonated gas phase dimer, the benchmark quantum trajectory was run for 125 ps and each of the contracted simulations for at least 250 ps. For the deprotonated dimer, the benchmark trajectory was run for 120 ps and each contracted simulation for at least 300 ps. For liquid water, the benchmark quantum trajectory was 80 ps and all other simulations were run for 100 ps. Uncontracted estimators were evaluated by post-processing 100 ps of the trajectory with configurations taken every 4 fs for the protonated and deprotonated dimers and every 10 fs for liquid water. Statistical error estimates on all the quantities reported were calculated using the bootstrapping method as 99 % confidence intervals[66].

IV. RESULTS

To assess the accuracy and efficiency of our AI-RPC scheme, we performed simulations of the protonated and deprotonated water dimers in the gas phase and of bulk liquid water. The former two were chosen for the reactive nature of the proton defects, which allow us to test the ability of our approach to capture on-the-fly making

and breaking of chemical bonds. Liquid water was chosen to demonstrate the applicability of our method to a condensed-phase system that exhibits a delicate balance between its NQEs. In particular, upon including NQEs water exhibits increased hydrogen bond deformation, which acts to weaken the hydrogen bonded network and destructure the liquid, and also increased proton sharing in the hydrogen bond, which has the opposite effect[42]. Liquid water thus allows us to assess if these competing quantum effects are correctly captured by AI-RPC.

A. Protonated water dimer

We first consider the protonated water dimer (an excess proton shared between two water molecules). In order to test the systematic convergence of our contraction scheme, Fig. 2 shows the quantum kinetic energy of the shared proton (top panel) and the dangling hydrogen nuclei (bottom panel) as the number of contracted replicas, P', is increased. Due to the protons' different chemical environments, the quantum kinetic energy of the weakly bound shared proton is 25 % less than that of the hydrogen nuclei, which are covalently bound to the water molecules. Already for a 1-replica contraction, which contracts the ring polymer to its centroid for the evaluation of the DFT interactions, both kinetic energies are within 3 % of the full quantum DFT result. This corresponds to an error of only 0.1 $k_\mathrm{B}T$ and is therefore already sufficiently converged for all practical purposes, and indeed is within the statistical sampling error of most AI-PIMD simulations. As the number of contracted replicas is increased, the result is observed to systematically converge to the exact quantum result by $P' = 8$, allowing high precision to be obtained by increasing the contraction level. In contrast, a standard PIMD simulation with 8 replicas gives an error in the kinetic energy of ∼23 %. Using uncontracted estimators (UE) to post-process the trajectories brings the error of the 1-replica contraction to below 0.05 $k_\mathrm{B}T$. In all cases, the oxygen nuclei kinetic energies were found to be within 0.01 $k_\mathrm{B}T$ of the full quantum DFT result, which is within the statistical error bars of our simulations. In addition, using MTS to increase the outer time step to 2 fs gave almost identical kinetic energies for $P' = 1$ and $P' = 6$ as well as for the UE for $P' = 1$ (not shown in the plot since the points overlap with those shown), compared to the 0.5 fs simulations, suggesting that the combination with MTS introduces negligible error.

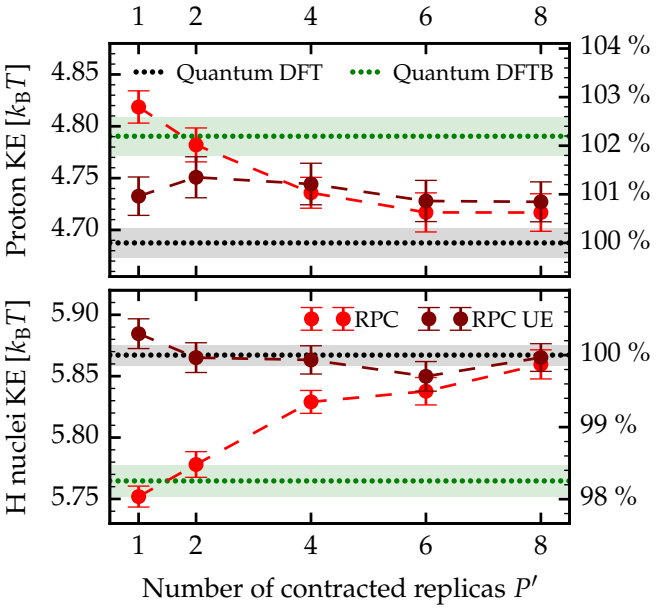

FIG. 2. Convergence of the proton kinetic energy with respect to the number of contracted replicas P' for the protonated water dimer. Kinetic energies from trajectories run with a fixed total number of replicas P=32 and an increasing number of contracted replicas P'=1 to P'=8 are shown in red. The values obtained using the uncontracted quantum kinetic energy estimator (UE) evaluated on geometries from the corresponding trajectories are shown in dark red. The top panel shows data for the dangling protons, while the bottom panel shows data for the shared proton. The quantum DFT and DFTB results, shown as dotted lines, were obtained from full 32-replica PIMD simulations. The axis on the right shows the percentage of the total kinetic energy obtained relative to the quantum DFT result. The classical contribution to the kinetic energy for each proton is 1.5 $k_\mathrm{B}T$. Statistical error estimates are indicated using error bars and shading.

To characterize the position of a proton between two oxygen atoms, we use the proton sharing coordinate δ defined as

$$\delta = d_\mathrm{OH} - d_\mathrm{O'H}, \tag{16}$$

where d_OH and $d_\mathrm{O'H}$ are the distances of the proton from the oxygen atoms. Fig. 3 shows the probability distribution along the δ coordinate. Here a contraction to the centroid combined with MTS using a 2 fs outer time step

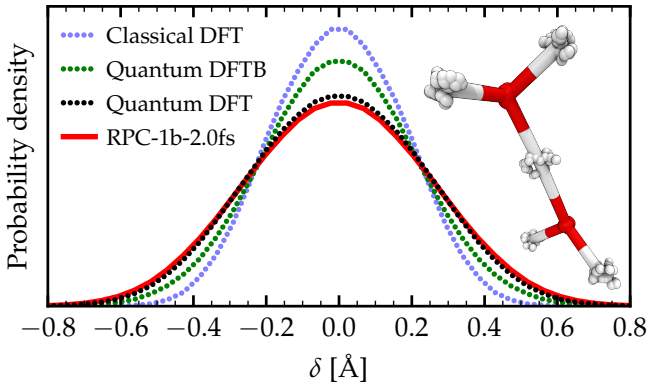

FIG. 3. Probability densities along the proton transfer coordinate, δ, for the protonated water dimer at 300 K. The quantum DFT and DFTB results were obtained from 32-replica PIMD simulations.

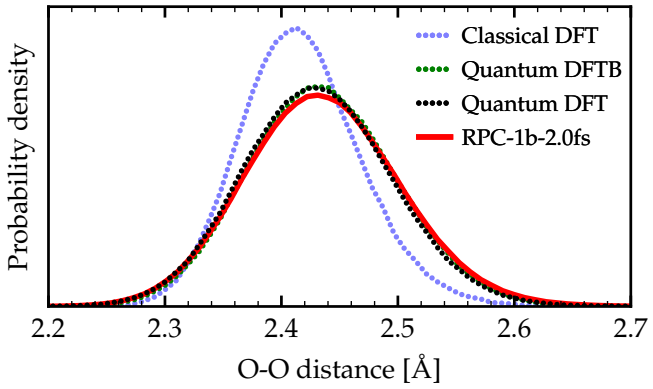

FIG. 4. Probability densities of the O-O distance for the protonated water dimer at 300 K. The quantum DFT and DFTB results were obtained from 32-replica PIMD simulations.

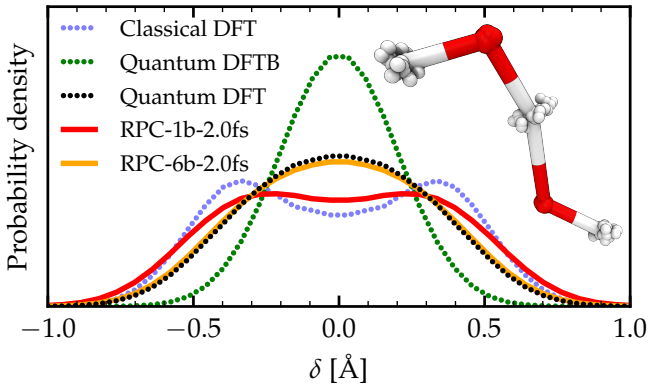

FIG. 5. Probability densities along the proton transfer coordinate, δ, for the deprotonated water dimer at 300 K. The quantum DFT and DFTB results were obtained from 32-replica PIMD simulations.

(RPC-1b-2.0fs) is sufficient to capture the broader distribution of proton positions within the statistical error bars, which is not captured by either classical DFT or the quantum simulation of the pure DFTB system. This wider distribution of the shared proton positions upon including NQEs changes the strength of the binding to the water molecules and thus also changes the O-O distance distribution shown in Fig. 4. Again, AI-RPC is in excellent agreement with the full quantum DFT results, even with a centroid contraction. The O-O distribution is also captured almost exactly by a pure quantum DFTB simulation. However, since the quantum DFTB simulation does not obtain the correct quantum proton δ distribution, this agreement is likely fortuitous.

B. Deprotonated water dimer

The deprotonated water dimer provides a particularly interesting test case, as NQEs cause a qualitative change in the behavior of the shared proton[67]. With classical nuclei, a double peak in the distribution of the proton sharing coordinate δ is observed, i.e. it exists primarily as a hydroxide ion bound to one of the water molecules, $H_2O{-}OH^-$, with frequent switching in which oxygen atom holds the proton defect. However, upon including NQEs this distribution changes to a single broad peak i.e. corresponding to a situation that more closely resembles a delocalized proton bound between two hydroxide like species, $[HO{-}H{-}OH]^-$. This presents a more challenging test case, since the DFTB reference system is a much poorer approximation to DFT for this problem. For example, in contrast to DFT, a classical DFTB simulation already predicts that the proton should be symmetrically tightly shared even without NQEs included.

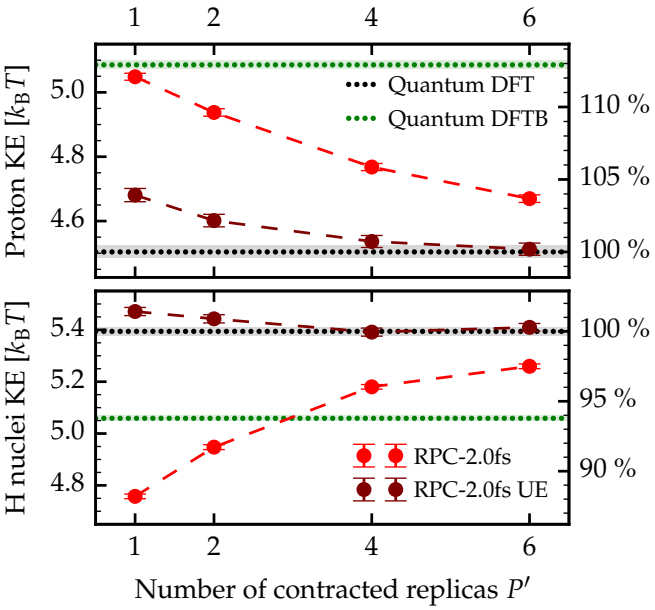

FIG. 6. Convergence of the proton kinetic energy with respect to the number of contracted beads P' for the deprotonated water dimer. Kinetic energies from trajectories run with a fixed total number of replicas P=32 and the number of contracted replicas from P'=1 to P'=6 and multiple time stepping (M=4, Δt=2.0 fs) are shown in red. The values obtained using the uncontracted quantum kinetic energy estimator (UE) evaluated on geometries from the corresponding trajectories are shown in dark red. The top panel shows data for the shared proton and the bottom panels shows the data for the dangling H nuclei. The quantum DFT and DFTB results, shown as dotted lines, were obtained from full 32-replica PIMD simulations. The axis on the right shows the percentage of the total kinetic energy obtained relative to the quantum DFT result. The classical contribution to the kinetic energy for each proton is 1.5 $k_B T$. The shading around the lines and error bars show the statistical error estimates.

The degradation in the reference system leads to a centroid contracted result (RPC-1b-2.0fs) for the proton δ distribution, shown in Fig. 5, which does not capture all of the quantum DFT result and still retains a remnant of the classical-like double peak structure. However, even with the poorer reference, increasing the number of contracted replicas to $P' = 6$ allows good agreement with the quantum DFT results to be obtained, again highlighting the systematic convergence of AI-RPC.

This slower convergence with the number of contracted replicas due to the poorer performance of the DFTB reference also manifests strongly in the quantum kinetic energies in Fig. 6, where a 1-replica contraction gives an error of 0.54 $k_\mathrm{B}T$ (11.9 %) for the shared proton. This error decreases to a much more acceptable value of 0.16 $k_\mathrm{B}T$ (3.5 %) when 6 contracted replicas are used, which still gives a considerable reduction in computational cost compared to a full quantum DFT calculation. In addition, using UEs reduces the error in the kinetic energy from a 1-replica contraction to within 0.17 $k_\mathrm{B}T$ (3.7 %) and 0.08 $k_\mathrm{B}T$ (1.5 %) of the exact result for the proton and H nuclei, respectively. Using UEs with 4 contracted replicas, gives kinetic energies that are within the error bars of the full quantum DFT result.

C. Liquid water

Owing to the tight binding approximation in DFTB, the method generally performs better for covalently bonded systems, and shows notable deficiencies in describing the structure of weakly bound and hydrogen bonded systems[68–70]. This problem is apparent for liquid water, where, as shown in top panel of Fig. 7, the quantum DFTB O-O radial distribution function (RDF) lacks any structure beyond the first peak[71]. From this, one might expect that an AI-RPC scheme based on DFTB would converge very slowly for liquid water as the number of contracted replicas is increased. However, as shown in Fig. 7, AI-RPC with contraction to a single replica yields RDFs that are graphically identical to full quantum DFT, thus further demonstrating the robustness of our approach.

The reason for the rapid convergence can be elucidated by examining the O-H and H-H RDFs obtained from pure DFTB, where the first (intramolecular) peaks show good agreement, while the remaining ones are noticeably shifted and broadened compared to DFT. As highlighted in Sec. II B, the main requirement for a good choice of the reference force for our AI-RPC scheme is that it yields a smoothly varying difference force when subtracted from the full DFT force. Hence, while DFTB does not capture the smoothly varying interactions at longer ranges in DFT correctly, these are corrected by the full DFT interactions, evaluated on the contracted replicas.

The accuracy of AI-RPC using a DFTB reference force is also apparent in the rapid convergence of the local geometries of the hydrogen bonds in water engendered in

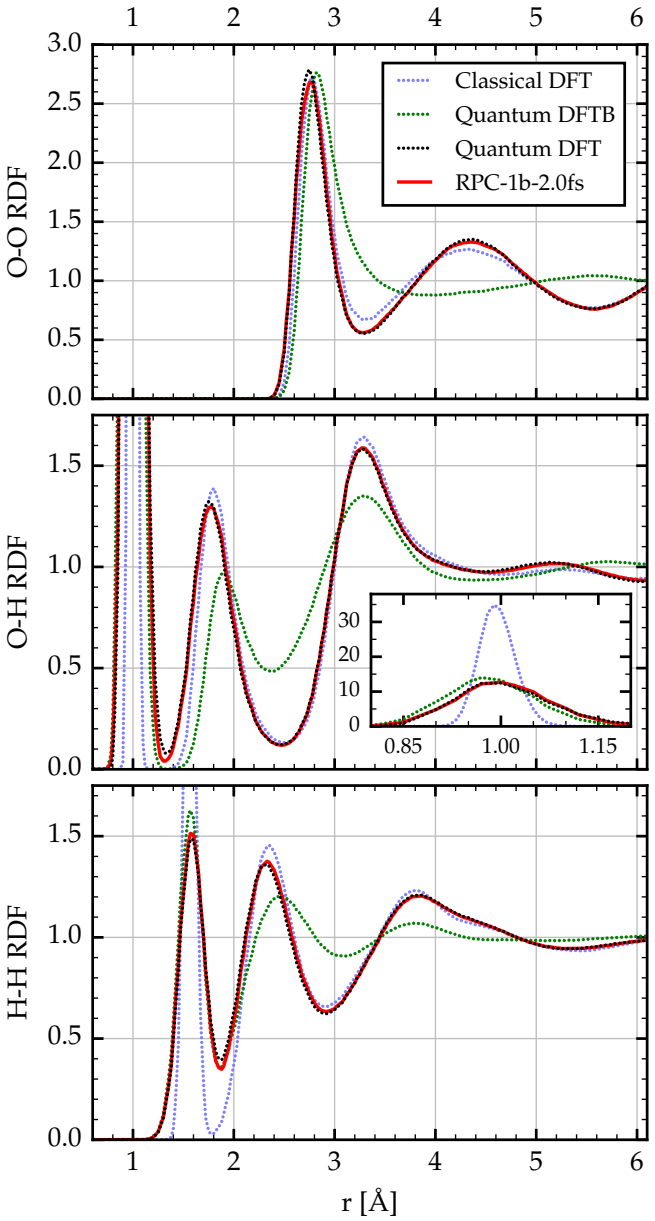

FIG. 7. Radial distribution functions of liquid water at 300 K obtained from AI-RPC and comparison simulations. The quantum DFT results were obtained from a 6-replica PIGLET simulation, while the quantum DFTB results were obtained using a 32-replica PIMD simulation. A 1-replica contraction with a 2 fs outer time step (RPC-1b-2.0fs) is observed to give results graphically identical to the full quantum DFT result.

the distribution of the proton sharing coordinate δ in Fig. 8 and in the distribution of hydrogen bond angles θ in Fig. 9. These distributions, which show large quantum effects, describe the delocalization of protons in hydrogen bonds in the direction along the hydrogen bond (δ) and perpendicular to it (θ). The quantum effects on these two coordinates are in competition, with the shift towards more shared protons (values of δ closer to 0) strength-

ening the hydrogen bond network, and the larger range of angles explored (values of θ further from 180°) weakening it. Therefore, the ability to describe both of these distributions accurately is closely related to the ability to describe the delicate balance of the competing quantum effects in liquid water. All the geometric properties obtained from our 1-replica contraction simulation were identical to those from simulations with contraction to 6 replicas, confirming that centroid contraction has already converged to the result that would be obtained from a full 32 bead PIMD simulation.

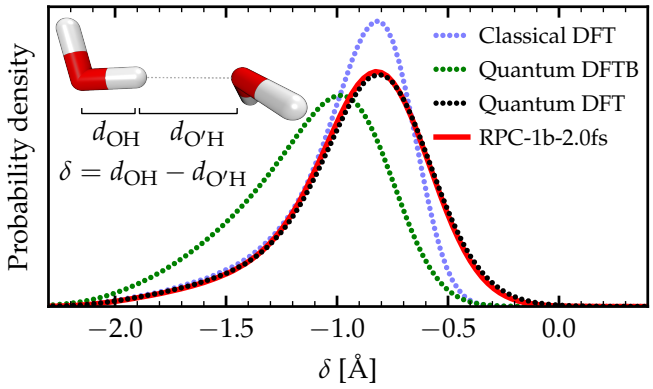

FIG. 8. Distribution of the proton transfer coordinate, δ, in liquid water at 300 K from AI-RPC and comparison simulations. The quantum DFT results shown obtained from a 6-replica PIGLET simulation, while the quantum DFTB results were obtained using a 32-replica PIMD simulation. A 1-replica contraction with a 2 fs outer time step (RPC-1b-2.0fs) is observed to give results graphically identical to the full quantum DFT result. The atom distances are labeled in a snapshot of a hydrogen bond.

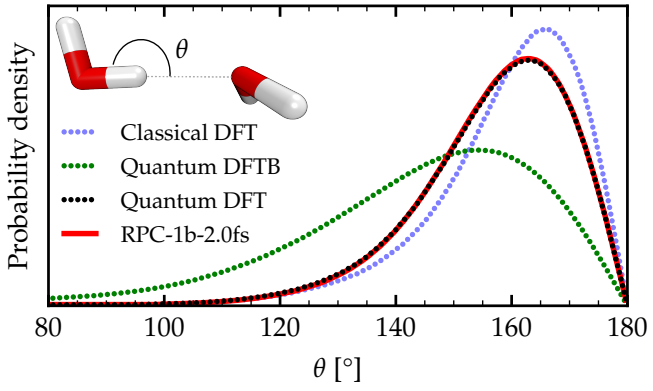

FIG. 9. Distribution of the hydrogen bond angle, θ, in liquid water at 300 K from AI-RPC and comparison simulations. The quantum DFT results shown obtained from a 6-replica PIGLET simulation, while the quantum DFTB results were obtained using a 32-replica PIMD simulation. A 1-replica contraction with a 2 fs outer time step (RPC-1b-2.0fs) is observed to give results graphically identical to the full quantum DFT result. The angle θ is labeled in a snapshot of a hydrogen bond.

Competing quantum effects are also known to manifest in the quantum kinetic energy of liquid water. Hence, finally we consider the convergence of the quantum kinetic energy of H and O nuclei in liquid water, shown in Fig. 10. Again, systematic convergence to the quantum DFT result is obtained with an error even with a contraction to 1 replica being within 1.2 %, or 0.07 k_BT of the exact result for the H nuclei, which becomes indistinguishable from the statistical error bars of the exact result (below 0.05 %) with a 6-replica contraction. Using UEs on the centroid contracted trajectory reduces the differences in the quantum kinetic energy from the benchmark to 0.74 % for H nuclei and 0.04 % for O nuclei, within the statistical error bar of the exact result of our 100 ps simulations.

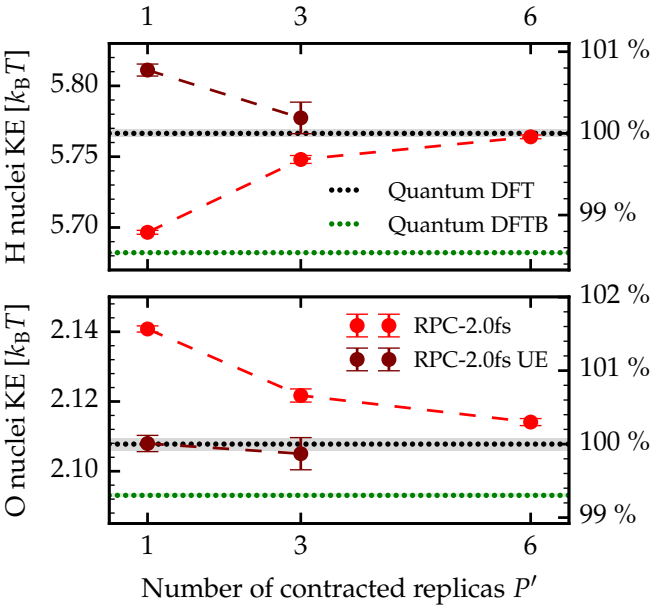

FIG. 10. Convergence of the hydrogen (top panel) and oxygen (bottom panel) nuclei quantum kinetic energy in liquid water with respect to the number of contracted replicas P'. Kinetic energies from trajectories run with a fixed total number of replicas P=32 and an increasing number of contracted replicas P'=1,3,6 and multiple time stepping (M=4, Δt=2.0 fs) are shown in red. The values obtained using the uncontracted quantum kinetic energy estimator (UE) evaluated on geometries from the P'=1 and P'=3 trajectories are shown in dark red. The quantum DFT results were obtained from a 6-replica PIGLET simulation, while the quantum DFTB results were obtained using a full 32-replica PIMD simulation. The axis on the right shows the percentage of the total kinetic energy obtained relative to the quantum DFT result. The classical contribution to the kinetic energy for each nucleus is 1.5 k_BT. Statistical error estimates are indicated using error bars and shading.

V. DISCUSSION AND CONCLUSIONS

In summary, our results show that for both liquid water and the reactive protonated and deprotonated dimer systems, our AI-RPC scheme for DFT, using an SCC-DFTB reference force, is able to obtain quantitative accuracy for a wide range of structural and energetic properties using contraction to the centroid. Indeed, to distinguish the 1-replica (centroid) contraction results from the statistical error bars of the exact quantum DFT results required trajectories of over 100 ps, which are much longer than those typically performed for AIMD simulations (in total in this study we performed over 3 ns of PIMD and AI-RPC for the gas phase systems and over 300 ps for the liquid water system). When compared to the impact of other subtle factors which can affect the results of AIMD simulations, such as basis set choice[72] and finite system size effects, such small discrepancies are likely to be unnoticeable in practical applications.

A centroid contracted AI-RPC simulation provides considerable computational speed-ups. Combined with an MTS scheme, a 1-replica contraction simulation reduces the number of DFT electronic structure calculations to evolve 2 fs in time from 128 (32 replicas and 4 time steps) to 1. This speed-up of two orders of magnitude would only be attainable if the cost of DFTB was truly negligible compared to DFT. In practice, running the current implementation on a single 16-core node (2x Intel Xeon E5-2670), our centroid contracted simulation of the 64 molecule liquid water system gave 5.6 ps of dynamics per day, which is exactly the same as we obtained from a standard classical AIMD simulation (without MTS). This represents a 35× speed-up over a full 32-replica DFT simulation on the same hardware. Even greater performance gains can be obtained if the difference in the computational cost of the two electronic structure methods was increased. This would be the case with a faster DFTB implementation or when the DFT calculations become more expensive, for example with a larger basis set. Additionally, the protonated and deprotonated dimer systems demonstrated that DFTB can act as a good reference potential for the hybrid B3LYP functional, where due to the higher computational cost of hybrid functionals, the speed-ups for condensed phase systems would be nearer the theoretical maximum. The reason our AI-RPC approach fared well for both a hybrid (B3LYP) and standard GGA (revPBE) exchange correlation functional is that, as explained in Sec. II B, the reference potential simply needs to capture the broad essence of the full potential to leave a smooth difference potential. Indeed, since the difference of SCC-DFTB from any given density functional is so large compared to the variation among them, it should be an equally good reference for most functionals, making this combination widely applicable. However, ongoing improvements to DFTB's ability to treat intermolecular interactions [34,71] offer the opportunity to bring it closer to DFT methods and allow even larger time steps and more aggressive contractions to be achieved.

If even tighter accuracy is required, our results demonstrate that AI-RPC exhibits systematic convergence with the number of contracted beads and that uncontracted estimators provide a way to increase the accuracy of the results by post-processing. Also, by examining the difference between the uncontracted and contracted estimator, one can assess whether convergence has been obtained.

To understand the success of AI-RPC, it is important to note that there is a fundamental difference between the performance of the pure reference potential and a centroid contraction. In AI-RPC the centroid positions are still determined entirely by the full DFT forces and only the higher normal modes of the ring polymers are coupled to the reference forces. Hence it is instructive to consider the most extreme case, where the reference forces are zero. This case corresponds to moving under the exact centroid force and then averaging the observables over the free ring polymer distribution (averaged over the free particle position uncertainty). For interacting systems, this free ring polymer distribution can be noticeably perturbed, e.g. in liquid water it is much more confined along the O-H stretch coordinate. However, the perturbation to the distribution is dominated by very strong forces in the physical potential, such as chemical bonds. Therefore, the change in the shape of the ring polymer can be accurately obtained using a reference potential which is able to broadly capture these strong forces. The local ring polymer distribution produced by the reference potential can then be refined using additional full force evaluations by increasing the contraction level.

Our results also suggest several potential directions for future development. In particular, the success of the uncontracted estimators, which can be recognized to arise as the zeroth-order cumulant in the expansion of the reweighted trajectory[22], indicates that reweighting the configurations generated from contracted simulations to the full uncontracted Hamiltonian could provide an even larger benefit to properties such as the kinetic energy and also allow position based properties to be improved. In a similar spirit, due to the well defined Hamiltonian, AI-RPC could also enhance recently proposed path integral replica exchange schemes[73] by allowing exchanges between different contractions levels, removing the current need to change the number of degrees of freedom upon exchange. If even higher accuracy is required, one could also envisage using AI-RPC to generate trial moves in a hybrid Monte Carlo scheme[74]. Finally, analogously to MTS, RPC is not limited to two levels of contraction and hence could be combined with an even higher tier *ab inito* method[50] at a third level to achieve very high accuracy at moderate cost. For example, a three-level contraction scheme could combine DFTB on all the replicas, DFT on an intermediate number of replicas and MP2 or another correlated method on the centroid.

In conclusion, our AI-RPC approach provides a method to routinely include nuclear quantum effects ex-

actly in the static equilibrium properties of *ab initio* molecular dynamics simulations. Additionally, this development offers a unique way to reduce the computational cost of *ab initio* CMD and RPMD dynamics, enabling previously unfeasible applications.

ACKNOWLEDGMENTS

The authors thank Will Pfalzgraff and Lu Wang for helpful comments and a thorough reading of this manuscript. We also thank Lu Wang for providing the revPBE-D3 water quantum DFT trajectory. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under Award Number DE-SC0014437. T.E.M also acknowledges support from a Cottrell Scholarship from the Research Corporation for Science Advancement and an Alfred P. Sloan Research fellowship.

[1] M. Parrinello and A. Rahman, J. Chem. Phys. **80**, 860 (1984).
[2] D. Chandler and P. G. Wolynes, J. Chem. Phys. **74**, 4078 (1981).
[3] R. P. Feynman and A. R. Hibbs, *Quantum mechanics and path integrals* (McGraw-Hill New York, 1965).
[4] D. Marx and M. Parrinello, Zeitschrift fur Phys. B Condens. Matter **95**, 143 (1994).
[5] D. Marx and M. Parrinello, J. Chem. Phys. **104**, 4077 (1996).
[6] M. E. Tuckerman, D. Marx, M. L. Klein, and M. Parrinello, J. Chem. Phys. **104**, 5579 (1996).
[7] J. Liu, W. H. Miller, F. Paesani, W. Zhang, and D. A. Case, J. Chem. Phys. **131**, 164509 (2009).
[8] S. Habershon, B. J. Braams, and D. E. Manolopoulos, J. Chem. Phys. **127**, 174108 (2007).
[9] E. Rabani, D. R. Reichman, G. Krilov, and B. J. Berne, Proc. Natl. Acad. Sci. U. S. A. **99**, 1129 (2002).
[10] I. R. Craig and D. E. Manolopoulos, J. Chem. Phys. **121**, 3368 (2004).
[11] S. Habershon, D. E. Manolopoulos, T. E. Markland, and T. F. Miller, Annu. Rev. Phys. Chem. **64**, 387 (2013).
[12] J. Cao and G. A. Voth, J. Chem. Phys. **101**, 6184 (1994).
[13] R. W. Hall and B. J. Berne, J. Chem. Phys. **81**, 3641 (1984).
[14] E. Pollock and D. Ceperley, Phys. Rev. B **30**, 2555 (1984).
[15] M. E. Tuckerman, B. J. Berne, G. J. Martyna, and M. L. Klein, J. Chem. Phys. **99**, 2796 (1993).
[16] M. Ceriotti, M. Parrinello, T. E. Markland, and D. E. Manolopoulos, J. Chem. Phys. **133**, 124104 (2010).
[17] M. Ceriotti, D. E. Manolopoulos, and M. Parrinello, J. Chem. Phys. **134**, 084104 (2011).
[18] M. Ceriotti and D. E. Manolopoulos, Phys. Rev. Lett. **109**, 100604 (2012).
[19] M. Shiga and W. Shinoda, J. Chem. Phys. **123**, 134502 (2005).
[20] W. Shinoda and M. Shiga, Phys. Rev. E **71**, 041204 (2005).
[21] A. Pérez and M. E. Tuckerman, J. Chem. Phys. **135**, 064104 (2011).
[22] M. Ceriotti, G. A. R. Brain, O. Riordan, and D. E. Manolopoulos, Proc. R. Soc. A Math. Phys. Eng. Sci. **468**, 2 (2011), arXiv:1107.1908.
[23] L. Wang, M. Ceriotti, and T. E. Markland, J. Chem. Phys. **141**, 104502 (2014).
[24] M. Ceriotti, J. Cuny, M. Parrinello, and D. E. Manolopoulos, Proc. Natl. Acad. Sci. U. S. A. **110**, 15591 (2013).
[25] L. Wang, S. D. Fried, S. G. Boxer, and T. E. Markland, Proc. Natl. Acad. Sci. **111**, 18454 (2014).
[26] "GLE Input generator, https://epfl-cosmo.github.io/gle4md/index.html?page=matrix," .
[27] M. Ceriotti and T. E. Markland, J. Chem. Phys. **138**, 014112 (2013).
[28] T. E. Markland and D. E. Manolopoulos, J. Chem. Phys. **129**, 024105 (2008).
[29] T. E. Markland and D. E. Manolopoulos, Chem. Phys. Lett. **464**, 256 (2008).
[30] G. S. Fanourgakis, T. E. Markland, and D. E. Manolopoulos, J. Chem. Phys. **131**, 94102 (2009).
[31] T. Frauenheim, G. Seifert, M. Elstner, Z. Hajnal, G. Jungnikel, D. Porezag, S. Suhai, and R. Scholz, Phys. Stat. Sol. **217**, 41 (2000).
[32] M. F. Herman, E. J. Bruskin, and B. J. Berne, J. Chem. Phys. **76**, 5150 (1982).
[33] M. Tuckerman, B. J. Berne, and G. J. Martyna, J. Chem. Phys. **97**, 1990 (1992).
[34] M. Gaus, Q. Cui, and M. Elstner, J. Chem. Theory Comput. **7**, 931 (2011).
[35] E. Anglada, J. Junquera, and J. Soler, Phys. Rev. E **68**, 55701 (2003).
[36] M. Guidon, F. Schiffmann, J. Hutter, and J. VandeVondele, J. Chem. Phys. **128**, 1 (2008).
[37] N. Luehr, T. E. Markland, and T. J. Martínez, J. Chem. Phys. **140**, 084116 (2014).
[38] H. Y. Geng, J. Comput. Phys. **283**, 299 (2015).
[39] R. P. Steele, J. Chem. Phys. **139**, 011102 (2013).
[40] S. Fatehi and R. P. Steele, J. Chem. Theory Comput. **11**, 884 (2015).
[41] M. Del Ben, J. Hutter, and J. VandeVondele, J. Chem. Phys. **143**, 054506 (2015).
[42] S. Habershon, T. E. Markland, and D. E. Manolopoulos, J. Chem. Phys. **131**, 024501 (2009).
[43] T. E. Markland, J. A. Morrone, B. J. Berne, K. Miyazaki, E. Rabani, and D. R. Reichman, Nat. Phys. **7**, 134 (2011).
[44] T. E. Markland, J. A. Morrone, K. Miyazaki, B. J. Berne, D. R. Reichman, and E. Rabani, J. Chem. Phys. **136**, 074511 (2012), arXiv:1111.4191.
[45] T. E. Markland and B. J. Berne, Proc. Natl. Acad. Sci. U. S. A. **109**, 7988 (2012).
[46] G. Han, Y. Deng, J. Glimm, and G. Martyna, Comput. Phys. Commun. **176**, 271 (2007).
[47] T. D. Hone, P. J. Rossky, and G. A. Voth, J. Chem. Phys. **124**, 154103 (2006).
[48] J. A. Morrone, T. E. Markland, M. Ceriotti, and B. J. Berne, J. Chem. Phys. **134**, 014103 (2011).
[49] M. Rossi, M. Ceriotti, and D. E. Manolopoulos, J. Chem. Phys. **140**, 234116 (2014), arXiv:1406.1074v1.
[50] V. Kapil, J. VandeVondele, and M. Ceriotti, J. Chem. Phys. **144**, 054111 (2016), arXiv:1512.00176.
[51] G. Bussi, D. Donadio, and M. Parrinello, J. Chem. Phys. **126**, 014101 (2007).
[52] M. Ceriotti, J. More, and D. E. Manolopoulos, Comput. Phys. Commun. **185**, 1019 (2013).
[53] A. D. Becke, J. Chem. Phys. **98**, 5648 (1993).
[54] M. J. Frisch, G. W. Trucks, H. B. Schlegel, G. E. Scuseria, M. A. Robb, J. R. Cheeseman, G. Scalmani, V. Barone, B. Mennucci, G. A. Petersson, H. Nakatsuji, M. Caricato, X. Li, H. P. Hratchian, A. F. Izmaylov, J. Bloino, G. Zheng, J. L. Sonnenberg, M. Hada, M. Ehara, K. Toyota, R. Fukuda, J. Hasegawa, M. Ishida, T. Nakajima, Y. Honda, O. Kitao, H. Nakai, T. Vreven, J. A. Montgomery, J. E. Peralta, F. Ogliaro, M. Bearpark, J. J. Heyd, E. Brothers, K. N. Kudin, V. N. Staroverov, R. Kobayashi, J. Normand, K. Raghavachari, A. Rendell, J. C. Burant, S. S. Iyengar, J. Tomasi, M. Cossi, N. Rega, J. M. Millam, M. Klene, J. E. Knox, J. B. Cross, V. Bakken, C. Adamo, J. Jaramillo, R. Gomperts, R. E. Stratmann, O. Yazyev, A. J. Austin, R. Cammi, C. Pomelli, J. W. Ochterski, R. L. Martin, K. Morokuma, V. G. Zakrzewski, G. A. Voth, P. Salvador, J. J. Dannenberg, S. Dapprich, A. D. Daniels, Farkas, J. B. Foresman, J. V. Ortiz, J. Cioslowski, and D. J. Fox, "Gaussian 09, Revision C.01," (2009).
[55] S. Bahn and K. Jacobsen, Comput. Sci. Eng. **4**, 56 (2002).

[56] J. Hutter, M. Iannuzzi, F. Schiffmann, and J. VandeVondele, Wiley Interdiscip. Rev. Comput. Mol. Sci. **4**, 15 (2014).
[57] J. VandeVondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing, and J. Hutter, Comput. Phys. Commun. **167**, 103 (2005).
[58] J. Perdew, K. Burke, and M. Ernzerhof, Phys. Rev. Lett. **77**, 3865 (1996).
[59] Y. Zhang and W. Yang, Phys. Rev. Lett. **80**, 890 (1998).
[60] S. Grimme, J. Antony, S. Ehrlich, and H. Krieg, J. Chem. Phys. **132**, 154104 (2010).
[61] S. Goedecker, M. Teter, and J. Hutter, Phys. Rev. B **54**, 1703 (1996).
[62] G. Lippert, J. Hutter, and M. Parrinello, Mol. Phys. **92**, 477 (1997).
[63] J. VandeVondele and J. Hutter, J. Chem. Phys. **118**, 4365 (2003).
[64] J. Kolafa, J. Comput. Chem. **25**, 335 (2004).
[65] D. Richters and T. D. Kühne, J. Chem. Phys. **140**, 134109 (2014).
[66] P. Young, *Everything You Wanted to Know About Data Analysis and Fitting but Were Afraid to Ask*, SpringerBriefs in Physics (Springer International Publishing, 2015) arXiv:1210.3781.
[67] M. E. Tuckerman, D. Marx, M. L. Klein, and M. Parrinello, Science **275**, 817 (1997).
[68] P. Koskinen and V. Mäkinen, Comput. Mater. Sci. **47**, 237 (2009), arXiv:0910.5861.
[69] C. M. Maupin, B. Aradi, and G. A. Voth, J. Phys. Chem. B **114**, 6922 (2010).
[70] P. Goyal, H.-J. Qian, S. Irle, X. Lu, D. Roston, T. Mori, M. Elstner, and Q. Cui, J. Phys. Chem. B **118**, 11007 (2014).
[71] H. Hu, Z. Lu, M. Elstner, J. Hermans, and W. Yang, J. Phys. Chem. A **111**, 5685 (2007).
[72] Z. Ma, Y. Zhang, and M. E. Tuckerman, J. Chem. Phys. **137**, 044506 (2012).
[73] Y. Peng, Z. Cao, R. Zhou, and G. A. Voth, J. Chem. Theory Comput. **10**, 3634 (2014).
[74] R. Iftimie and J. Schofield, J. Chem. Phys. **114**, 6763 (2001).